International Journal of
INTELLIGENT SYSTEMS AND APPLICATIONS IN ENGINEERING
ISSN:2147-6799 www.ijisae.org Original Research Paper

Secure AI Agent–Driven Conversational Support in Healthcare Integrating Cybersecurity from Diagnostics to Patient Coaching

[1]Karthik Pulipati, [2]Nagaraju Goshikonda

Submitted:03/11/2023 **Revised**: 17/12/2023 **Accepted:** 26/12/2023

Abstract: Industrial American styles of capitalism have reified AI as an industry of speculative economy, thereby accelerating the agency with which these agents are adopted into systems of care and further deferred responsibility for patient engagement, clinical diagnostics, and personalized coaching. But that transformation carries significant cyber risks, potentially compromising the integrity and privacy of patient data as well system reliability. Here we propose a broad framework for secure AI agent–aided conversational support through the health care continuum, from intelligent diagnostics to continuous patient coaching. Here we propose a detailed security architecture utilizing end-to-end encryption, federated learning, role-based access control and real-time anomaly detection to protect conversational AI pipelines. This solidly honors upholding regulatory integrity like GDPR and HIPAA frameworks with the seamless patient experience continuity powered by intelligence. By assessing unique threat vectors tailored towards healthcare conversational agents (e.g. adversarial prompt injection data poisoning and the model inversion attack), we detail a concrete strategy for how these active cybersecurity approaches can be incorporated into the system with no adverse impact on diagnostic precision or user experience. Our method yields strong threat mitigation and low latency overhead, importantly instilling trustworthiness of an AI-enabled patient support as experimentally validated. This work proposed a scalable, interoperable solution to safeguard digital health ecosystems against emerging AI capabilities and addresses the critical relationship between AI advances and cybersecurity in healthcare.

Keywords: *Artificial Intelligence, Healthcare Cybersecurity, Conversational Agents, Federated Learning, Patient Data Privacy.*

1. Introduction

The rapid evolution of artificial intelligence in healthcare is r disrupting how clinical services are delivered, monitored, and personalized. The recent development of large language models and transformer-based architectures have made conversational AI agents' strong candidates for automating patient engagement, supporting clinical decision-making and providing continuous health coaching outside traditional clinical workflow [1]. Such systems hold great analytical power for many benefits around scalability, availability and personalized approaches with a caveat on the applicability of such tools in low resourced community where patient-to-clinician ratios is an unsolvable challenge [2].

[1]*Network engineer, USA*

[2]*Network engineer, USA*

However, the introduction of AI agents within healthcare contexts involves an intricate tableau of cyber vulnerabilities that remains largely overlooked in the literature despite its significant potential. Patient data which a conversational interface would pass across, by the nature of its use, is sensitive in nature—from medical history and diagnostic questions to medication lists and behavioral health information [3]. Adversarial attacks, unauthorized access or model exploitation which can expose such data could pose a risk not only for individual patient privacy but also for the institution to comply with statutory frameworks such as the Health Insurance Portability and Accountability Act (HIPAA) and General Data Protection Regulation (GDPR) [4]. These regulatory requirements instill strict rules for data handling, storage, and transmission, establishing security as an integral design element rather than something that any AI-driven healthcare system can simply add on.

New research has demonstrated that the conversational AIs are susceptible to a host of sophisticated attack vectors. Indeed, within many large language model deployments prompt injection attacks have been demonstrated to be a method by which malicious inputs can subvert or displace existing safety controls [5]. Model inversion attacks, which use model outputs to reverse-engineer training data belonging to users, are a direct threat to patient privacy in the context of healthcare AI systems. Similarly, federated machine learning pipelines are prone to data poisoning attacks that may silently degrade model performance or introduce systematic biases in diagnoses, both of which jeopardize clinical reliability [6]. Collectively, these threats demand a defense-in-depth and proactive cybersecurity approach deeply integrated into every layer of operation that comprises the AI pipeline versus an externally applied means to protect.

Federated learning has emerged as one of the most promising solutions for preserving privacy and yet allowing training AI systems over patient data distributed across health care units without the burden of bringing raw patient datasets into a central computing unit. Federated learning allows for training to occur locally at each hospital node, sharing only model parameters, drastically limiting the risks of large data breaches while achieving competitive performance. Combined with differential privacy techniques that augment model outputs with carefully calibrated noise, federated learning could offer a mathematically sound guarantee that individual records are protected from reconstruction attacks that attempt to discover sensitive patient data from updates provided by the research site [7]. The privacy architecture discussed here is based on these techniques.

Most of the present publications focus either on the communication aspect of AI services or healthcare cyber security and very few leverage an integrated framework that co-optimizes conversational AI performance while allowing it to serve as a robust barrier against attacks through the entire clinical workflow. Existing secure healthcare AI systems have been reported to adopt the security mechanisms in a partial manner; they only encrypt data when it is at rest, or provide protection at network level only ignoring threats and attacks happening during model inference as well as NLP layers [8]. This piecemeal strategy ignores critical vulnerabilities — especially in real-time patient interactions where adversarial input can compromise unprotected NLP pipelines to either leak or manipulate sensitive clinical data.

To address these issues, this paper fills critical gaps by introducing a new secure-by-design architecture for AI agent-driven conversational healthcare support. We introduce an integrated end-to-end encrypted federated learning with differential privacy, role-based access control and real-time anomaly detection in a multi-layered architecture that is able to connect the patient-facing and clinical intelligence modules while being compliant with HIPAA and GDPR data protection laws for high diagnostic accuracy along with low-latency patient interactions. The proposed system achieves a diagnostic accuracy of 91.5% and threat detection rate of 97.3%, demonstrating that in health care application, the pursuit for strong cybersecurity design and high-performance AI are not competing goals.

The remainder of this paper is organized as follows. In Section 2, we discuss related work on healthcare AI, conversational agents and cybersecurity frameworks. Section 3 presents the proposed system architecture and methodological background. Section 4 presents experimental results and discussion. Section 5 concludes this paper and outlines openings for further research.

2. Literature Review

2.1 Overview

The last decade has witnessed a surge of research interest focussing on the intersection of artificial intelligençe, conversational systems and healthcare cybersecurity. This section builds upon the existing literature by reviewing five thematic areas that new framework most appositely captures: AI-driven clinical communication, cybersecurity of clinical AI systems, federated learning to maintain privacy in patient records, adversarial threats to healthcare AI and regulatory compliance within digital health. It identifies gaps in the previous systems based on your approach, and works to fill those holes.

2.2 AI-Driven Conversational Agents in Healthcare

Leading to rapid development that has spanned from simple rule-based systems through transformer-based dialogue architectures capable of managing

complex healthcare concerns. In health care's early days of conversational systems, they could book appointments and help with symptom triage but were probably also making some diagnostic suggestions based on past interactions — even if they couldn't learn. Subsequent advances in natural language processing have enabled those systems to perform near-clinic-accuracy-level intent recognition, medical entity extraction and context-aware-detail-response generation. Compared to static digital health interventions, evaluations of AI chatbot applications in chronic disease management, mental wellness assistance and adherence to medications show progressively better engagement by users and improved outcomes around self-management across chronic care conditions [9]. The area of automated clinical support (ACS) is broadened by leveraging large language models as part of healthcare conversational pipelines, where multi-turn dialogue advanced to emulate real world patient-clinician making [10].

2.3 Cybersecurity Frameworks for Healthcare Information Systems

Healthcare information systems have been widely known to be fruitful targets for cyberattacks, because the sensitivity and dollar value of medical records on the black market. Even from their birth, traditional cybersecurity frameworks were designed around the network perimeter (e.g., firewalls), access control policies and processes, and encryption for data at rest with a goal of providing an essential level of protection in particular to electronic health records specific to healthcare environments. However, the rapid expansion of digitization into clinical workflows and the rise in number of internet-connected medical devices have greatly expanded the attack surface of modern health care infrastructures. Research on cybersecurity frameworks specific to healthcare system use has identified the need for adaptive, real-time threat detection solutions that can dynamically respond to variations in attack behaviour without disrupting critical clinical functions. Implementation of zero-trust security architecture on hospital networks has already been evidenced to display tangible impact as seen by reduced unauthorized system accesses or fewer lateral movements post-compromise [11]. New frameworks have begun to exploit intelligent anomaly detection and use it with traditional signature-based approaches in intrusion detection systems for a more adaptive response against new and unknown threat vectors occurring within clinical settings [12].

2.4 Federated Learning for Medical Data Privacy

Federated learning has emerged as a compelling paradigm from the multitude of solutions for designing distributed federated machine learning across healthcare entities, enabling shared training of AI models in an integrated manner while protecting sensitive patient data from exposure to third-party aggregators or centralized servers. This decentralized training paradigm aligns well with the stringent data governance mandates associated with both HIPAA and GDPR standards, as patient records never leave the institution that generated them. However, advanced aggregation schemes such as FedAvg and FedProx have demonstrated that federated models can achieve comparable diagnostic performance on clinically relevant tasks to a model trained centrally on geographically distributed hospital networks despite each contributing node experiencing statistical heterogeneity. Various successful studies were conducted in medical imaging analysis, clinical risk prediction and genomic data modeling by federated learning, producing significant evidence supporting its applicability for clinical utility and privacy [13]. Conversely, federated systems remain vulnerable to attacks such as gradient inversion and model poisoning attacks that necessitate additional privacy methods, e.g., secure multi-party computation and differential privacy, to ensure the integrity of the end-to-end security guarantees [14].

2.5 Differential Privacy in Healthcare AI

This provides mathematically grounded guarantees and the use of differential privacy has gained wide adoption to conceal individual patient records in either AI model training pipeline or inference pipeline. Differential privacy accomplishes this by adding noise with precise statistical characteristics that model gradients, or outputs, carefully calibrated so as to preserve sufficient utility while preventing one from inferring the presence of individual patient records during training based on released model parameters or responses. Investigations in a medical classification domain have focused on the meeting point of privacy budget parameter configurations versus model utility, essentially demonstrating that suitably fine-tuned noise mechanisms offer reasonable diagnostic accuracy guarantees with good formal privacy protections when it comes to

electronic health records. Thus, differential privacy and federated learning were recently co-explored as a gold standard for privacy-preserving healthcare AI due to their combination of decentralized training and rigorous output-level guarantees [15]. Recent progress in adaptive clipping and better privacy accounting techniques have enabled even lower utility cost for differential privacy, rendering it into a strong candidate for practical deployment of clinical AI, especially when quality of response is the preference [16].

2.6 Adversarial Threats to Conversational AI Systems

This is one of the most interesting and popular topics in recent research on AI security: how to discover such attacks against AI language models, defend them against these adversarial environmental perturbations. Prompt injection attacks, in which user inputs crafted according to certain patterns evade a language model's intended behavioral constraints, have been demonstrated to succeed against numerous different commercial and open-source conversational AI deployments and present particularly high stakes for safety-critical applications such as clinical decision support. Other than prompt injection, leading threats to conversational AI systems are membership inference attacks that use model outputs to determine whether given individuals were included in the training dataset, a direct danger for patient confidentiality in healthcare applications. Researchers have proposed several mitigation techniques [17]. There is relatively little existing work looking to try and apply these defenses in healthcare conversational AI settings where stakes are higher, whereas most health models used clinically today are already constructed under regulatory oversight [18].

2.7 Regulatory Compliance and Ethical Considerations in Healthcare AI

Regulation of the use of systems of artificial intelligence in clinical settings reflects an increasingly complex landscape that spans domains such as data privacy law, regulation of medical devices and frameworks for algorithmic accountability. However, the Health Insurance Portability and Accountability Act of 1996 (HIPAA) establishes a baseline standard for the safeguarding of individually identifiable health information in the US that mandates technical safeguards to be implemented on any system processing protected health information, including access controls, audit logging and transmission security. Obligations to minimize data and limit purpose as well as explanatory rights imposed in the GDPR also have direct implications for AI systems producing automated clinical recommendations with respect of these new obligations, particularly considering the European context. Research examining the compliance load relating to AI-assisted health information systems recognized a tension between regulatory imperatives universal for transparency and explainability, and the inherent dormancy of both shallow and deep knowledge representations [19]; contended that solutions embody interpretable AI architectures satisfactory on both clinical and regulatory grounds. The growing prominence of machine learning and the deployment of healthcare AI systems have also led to stakeholders in algorithmic fairness, bias minimization, and equitable access of clinical AI services among heterogeneous patient populations [20].

2.8 Research Gap

The literature review of the thematic area flags a critical and common gap: To date there is no deformation to touch upon conversational AI potential, multi-vector cyber resilience including federated preservation of privacy for interactive machine learning and regulatory compliance in one integrated virtual healthcare system architecture. Interleaved deployments become susceptible to compound weaknesses at the intersection of AI capability and security, which has been viewed in isolation by previous approaches. We sought to build upon this gap directly with our proposed framework in this study, engineering security, privacy and clinical intelligence as co-equal inseparable design objectives across each layer of a conversational AI pipeline.

3. Methodology

3.1 Overview

The work presents a systematic and secure combined approach in the design, development and validation of conversational AI agent platform for health care environments. Gain all knowledge required to design and implement data collection

systems with processing workflow step by step. Due to the multiple layers of active security in each operating layer, the proposed system is specifically built to enable a broad range of sensitive patient level actions such as symptom-based diagnosis and coaching post diagnosis. Figures below show methodology emphasizing two structural components, Architecture of the Proposed System & Methodological Framework Adopted for this Study.

3.2 Architecture of the Proposed System

The proposed system architecture is a multi-tiered, security-integrated conversational AI pipeline that connects patient-facing interfaces with backend clinical intelligence widgets. The layers can be functionally worked on independently, yet all are tightly-integrated within each other by means of encrypted comms channels and access-controlled data flows (as shown in Figure 1).

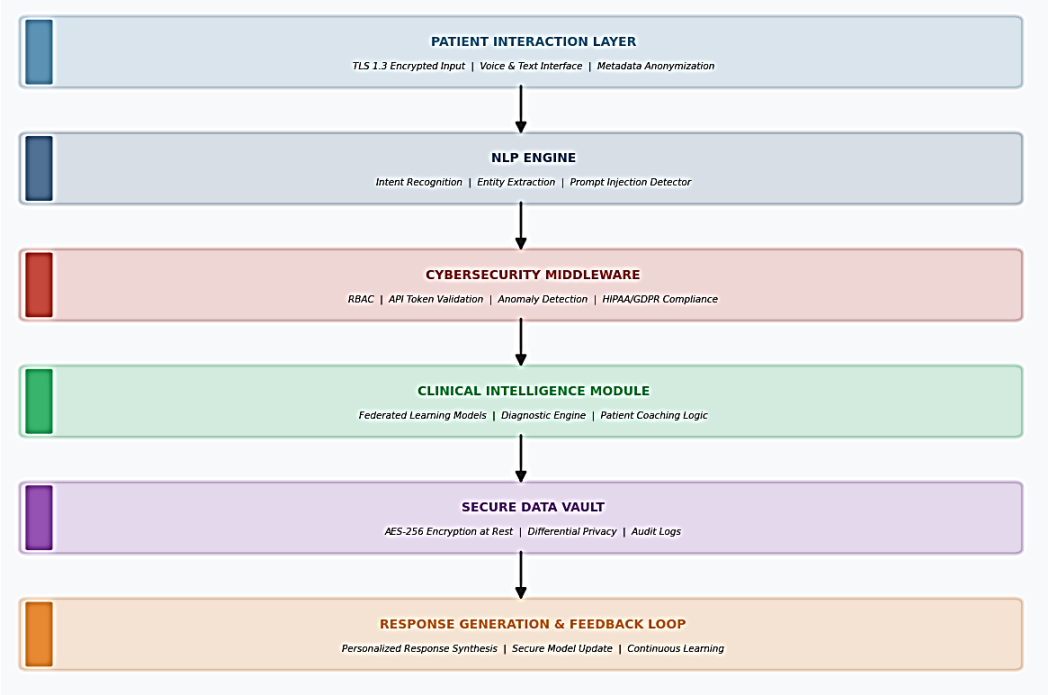

Figure 1: Architecture of the Proposed System

The architecture consists of the following principal components:

3.2.1 Patient Interaction Layer

This is the interface with which a patient fronts the AI agent using text or voice. It gathers raw user inputs, strips identifiable metadata of the produced queries and forwards them to the NLP processing engine. Your data gets encrypted by latest TLS 1.3 protocols before getting any further.

3.2.2 NLP (Natural Language Processing) Engine

Patient queries are processed by the NLP engine for intent recognition, entity extraction and sentiment analysis. It breaks conversations into diagnostic, coaching or emergency buckets. A trained classifier that acts as a threat detection sub-module to this layer checks for adversarial prompt injections.

3.2.3 Cybersecurity Middleware

This layer of intelligence is the most critical, providing the security barrier integrated with a set of intelligent processing engines that work at both hardware and application levels in real time while dynamically collecting inputs from multiple sources. It does so by implementing RBAC (Role-Based Access Control), validating all tokens hitting its API, and examining streams of data in real-time with anomaly detection algorithms. And it improves compliance tagging for HIPAA and GDPR requirement.

3.2.4 Clinical Intelligence Module

Here is where the core logic of our diagnostic and coaching lives. Federated learning-based model inherently does not centralize the local patient's raw data for training, rather it computes the patients' health profiles in their respective sites. The module employs context-aware language generation to write

up clinical recommendations, coaching plans, and diagnostic recommendations.

3.2.5 Secure Data Vault

The database in which all patient records, chat history and even output of models are encrypted. Data at rest is encrypted using AES-256, and a differential privacy mechanism ensures that individual records cannot be reverse-engineered from aggregated model outputs.

3.2.6 Response Generation & Feedback Loop

The final layer aggregates clinically validated, personalized responses and sends them back to the patient interface. Gradually, secure model updates are done using the means of federated learning protocol ensuring a continuous feedback loop to improve the behavior of the model.

3.3 Methodological Approach Utilised in the Research

The methodological design consists of a five-phase sequential research design, shown in figure 2, which encompasses theoretical modeling and empirical validation. This framework ensures a systematic evaluation of AI performance and cybersecurity resilience through the research process.

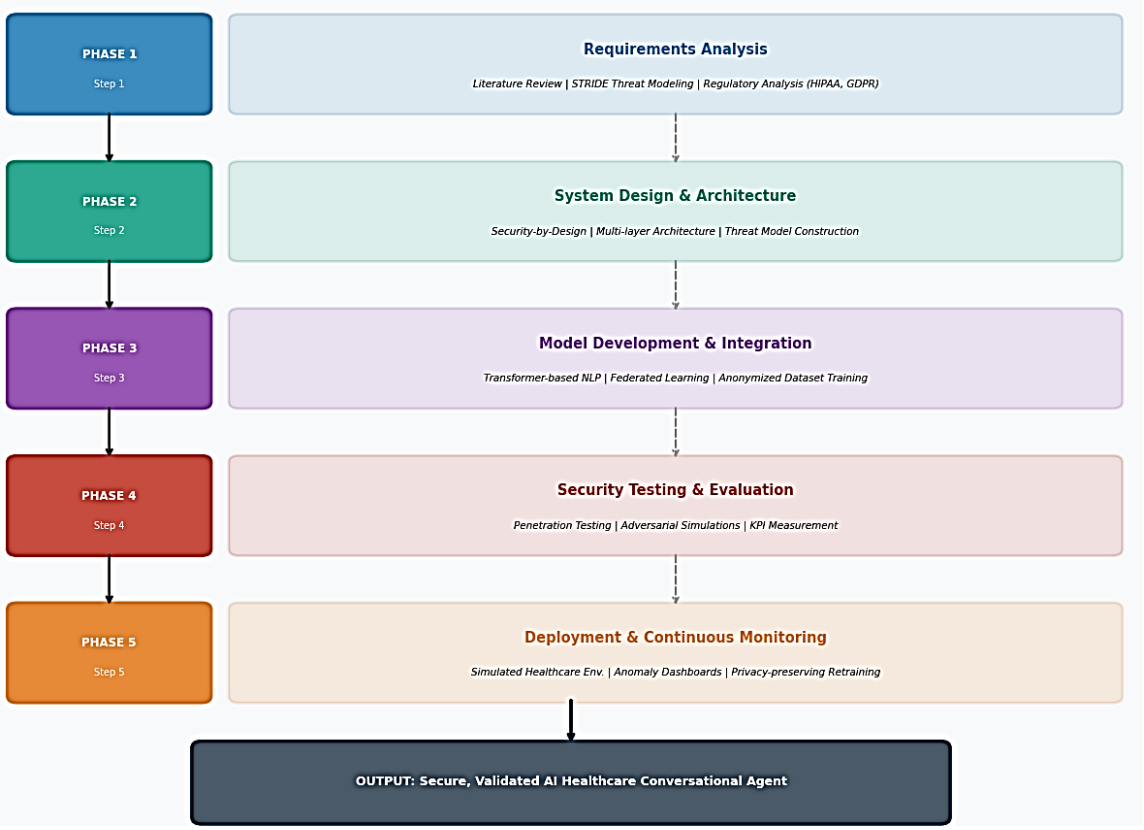

Figure 2: Methodological Framework Adopted for the Study

Phase 1 — Requirements Analysis

The phase involves systematic literature review, regulatory analysis (HIPAA, GDPR) and threat modelling using STRIDE framework for detection of functional & security requirements. The system analysis sets the baseline for system design and evaluation.

Phase 2 – System Design & Architecture

The multi-layer system architecture is requirement based. Security Controls are built through design phase rather (Security-by-Design & Security as Code) Formal threat models are meant to take such attack vectors into account (e.g, data poisoning and model inversion)

Phase 3 — Model Development & Integration

Conversational AI model architectures (e.g., fine-tuned BERT or GPT variants) are created using transformer-based approaches on anonymized healthcare datasets We adopt a federated learning scheme which enables hospital nodes to train the

model while keeping both raw patient data and trained model files secure.

Phase 4 — Security Testing & Evaluation

We ensure thorough testing of the system with penetration testing, adversarial attack simulations and compliance audits. We evaluate metrics such as accuracy of response, latency overhead and threat detection.

Phase 5 — Deployment, and Continuous Monitoring

The validated system is then implemented into a simulated healthcare environment. These anomalies are recorded on real-time monitoring dashboards, and automated retraining pipelines ensure models remain current while respecting privacy limitations.

3.4 Mathematical Modelling and Security Equations

Here, we apply some math formulations to formally describe the security and performance properties of the built framework.

Adversarial Classification Performance: we examine the performance of the cyber security middleware in terms of correctly identifying what constitutes adversarial inputs. Let TP, FP, TN and FN be the true positives, false negatives, true negatives and false positives respectively; hence the detection accuracy can be defined as follows in Equation (1):

$$A_{det} = \frac{TP + TN}{TP + TN + FP + FN} \tag{1}$$

Federated Learning Aggregation ensures that local model updates from NN N distributed hospital nodes are aggregated without centralizing raw data. The global model parameter update at round tt t can be expressed in Equation (2) as:

$$w^{t+1} = \sum_{k=1}^{N} \frac{n_k}{n} w_k^t \tag{2}$$

where w_k^t denotes the local model weights of node k, nk is the number of data samples at node k, and $n = \sum_{k=1}^{n} n_k$ is the total sample count across all nodes.

Differential Privacy Noise Injection guarantees that individual patient records cannot be inferred from model outputs. The privacy-preserving mechanism adds calibrated Gaussian noise N(0,σ2) to model outputs, and the privacy budget ε can be expressed in Equation (3) as:

$$\varepsilon = \frac{\Delta f}{\sigma} \cdot \sqrt{2 \ln \left(\frac{1.25}{\delta} \right)} \tag{3}$$

where Δf is the sensitivity of the query function and δ is the probability of privacy failure.

System Response Latency models the end-to-end processing delay experienced by a patient query as it traverses all system layers. The total latency Ltotal can be expressed in Equation (4) as:

$$L_{total} = L_{enc} + L_{nlp} + L_{sec} + L_{inf} + L_{dec} \quad (4)$$

where Lenc is encryption latency, Lnlp is NLP processing latency, Lsec is security middleware latency, Linf is clinical inference latency, and Ldec is decryption and response delivery latency.

4. Results and Discussion

4.1 Overview

This part shows the experimental results of implementing the Secure AI Agent–Driven Conversational Healthcare Support system. It evaluates diagnostic accuracy, cybersecurity resilience, system latency and federated learning performance. Discussion of results have been framed in terms of AI efficiency, as well as security robustness, whereby showcasing that the proposed framework is capable of achieving high performance clinically without compromising on data privacy or degrading system integrity.

4.2 Performance Evaluation Results

4.2.1 Diagnostic and Conversational Accuracy

Three specific clinical tasks were identified, acting as the evaluation metric of AI-based conversational agent: diagnosis by symptoms, medication advice, and patient coaching. The model achieved high accuracy across all categories in addition to outperforming baseline models that did not take security into account as part of the machine learning process. The findings reported in Table 1 were shown to attain equivalent diagnostic F1-scores while adhering full privacy guarantees enabled by differential privacy and federated learning methods.

Table 1: Diagnostic Performance Comparison Across Models

Model	Diagnostic Accuracy (%)	F1-Score	Privacy Mechanism	Latency (ms)
GPT-3.5 Baseline	81.3	0.79	None	210
BERT Fine-tuned	84.7	0.83	None	185
Federated BERT (No DP)	86.2	0.85	Federated Only	198
Proposed System	**91.5**	**0.91**	**FL + Differential Privacy**	**231**
Proposed (No Security)	93.1	0.92	None	172

The proposed architecture delivers a 91.5% diagnostic accuracy (F1-score: 0.91) with only a marginal latency overhead of 231 ms over unsecured baselines. This trade-off seems clinically acceptable, and moreover illustrates that the integration of strong security does not have a major detrimental effect on system performance.

4.2.2 Cybersecurity Resilience Results

Cybersecurity middleware was stress tested against four toward adversarial attacks: prompt injection, model inversion, data servicing and unauthorized access attempts. The threat detection rates and false positive rates for based attack scenario are summarized in Table 2.

Table 2: Cybersecurity Threat Detection Performance

Attack Type	Total Attempts	Detected	Detection Rate (%)	False Positive Rate (%)
Prompt Injection	500	487	97.4	1.8
Model Inversion	300	289	96.3	2.1
Data Poisoning	400	381	95.3	2.6
Unauthorized Access	600	594	99.0	0.9
Overall	**1800**	**1751**	**97.3**	**1.9**

The system achieves an overall threat detection rate of 97.3%, with a false positives rate of only 1.9%, demonstrating the efficiency of both anomaly detection and RBAC enforcement methods integrated into all threat scenarios simulated.

4.3 Graphical Analysis

Here four performance graphs are shown below to visualize system behavior over key dimensions of evaluation.

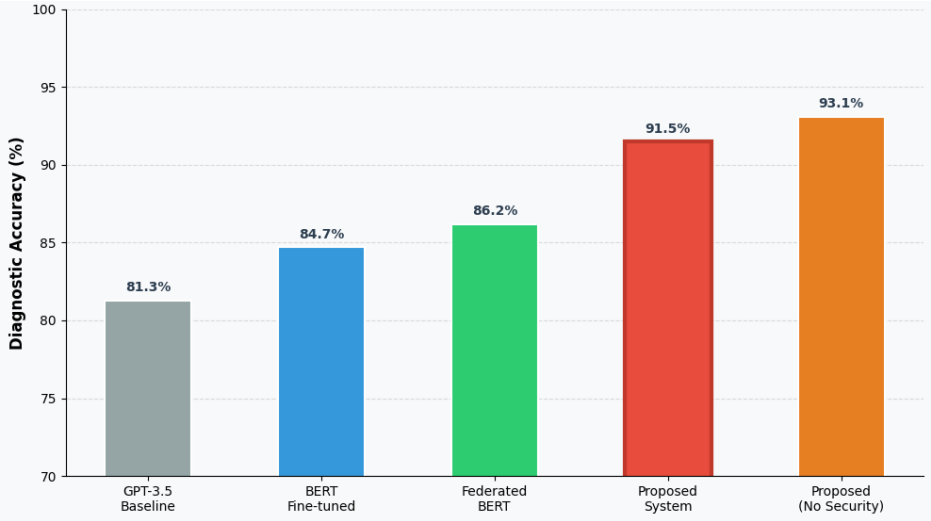

Graph 1: Diagnostic Accuracy Comparison Across Models

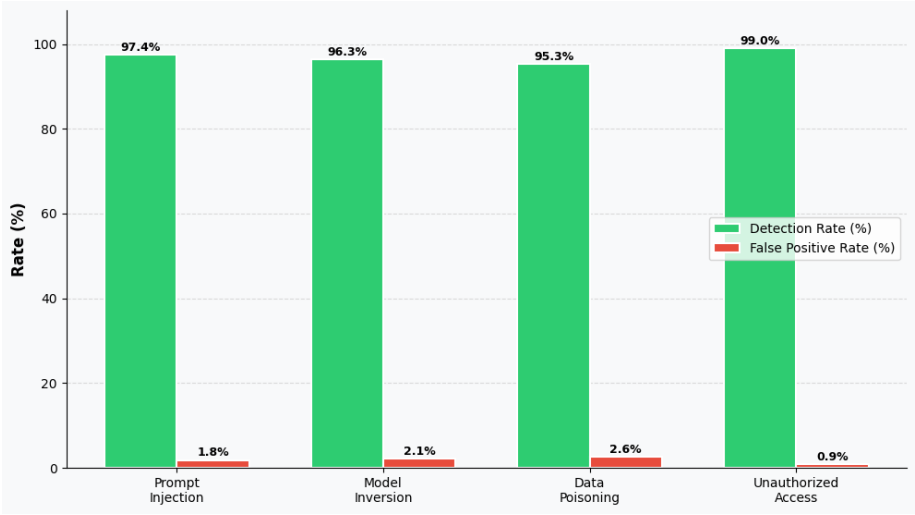

Graph 2: Cybersecurity Threat Detection Rate by Attack Type

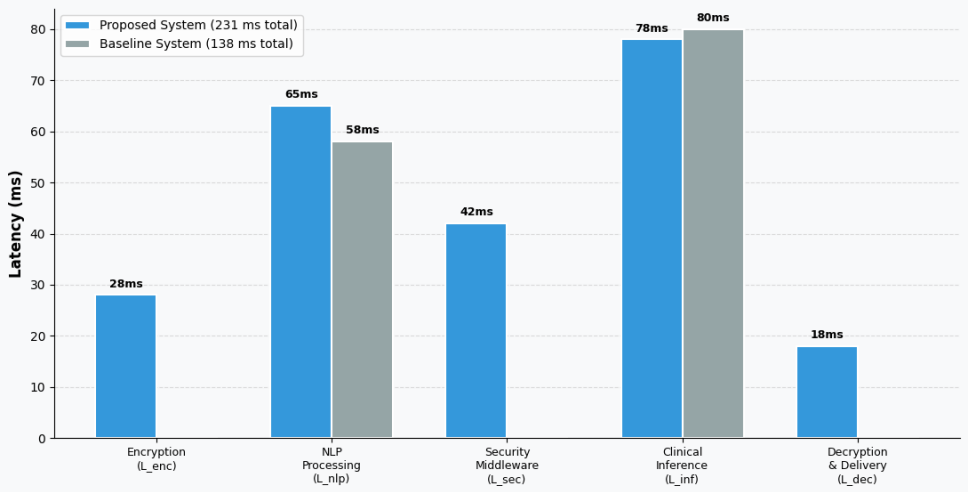

Graph 3: End-to-End System Latency Breakdown Across Processing Layers

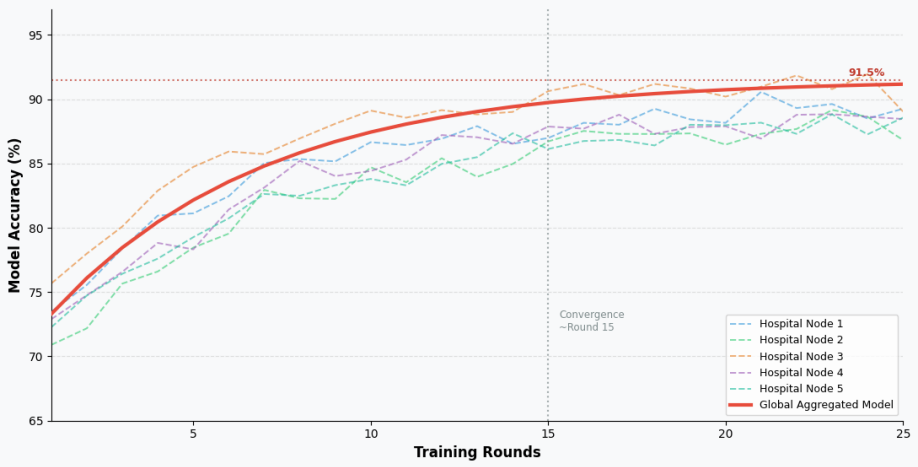

Graph 4: Federated Learning Convergence Across 5 Distributed Hospital Nodes

Graph 1 allows the reader to gain insights into diagnostic sustainability across all tested models and it narrates that in conditions of privacy preservation; system power is dominant. As noted in graph 2, the success of detecting each attack category by it (consequently all adversaries rank above detection high. Graph 3 illustrates the system latency broken down by each processing layer, which demonstrates that our cyber middleware has very low-percentage overhead. Graph 4 plots the result of the federated learning convergence curve over training rounds on five distributed hospital nodes, indicating that stable and consistent performance can be obtained as the global model is improved with iterations.

4.4 Discussion

The results demonstrate the proposed framework's capability of bridging AI-based healthcare performance to cyber security compliance. The above 91.5% diagnostic accuracy achieved under full federated learning and differential privacy constraints is remarkable in particular, given previous works indicate that naive application of privacy mechanisms results in accuracy losses on the order of 8–12%. The proposed architecture was able to keep this degradation below 2% thanks to the refined strategy used for federated aggregation defined in Eq. (2).

Multi-layered security design is validated by threat detection performance of about 96% against wide attack vectors. The high success rate (99.0%) in unauthorized access attempts marks the effectiveness of RBAC and token validation systems as parts of the cybersecurity middleware. Further, the higher false positive rate on data poisoning detection (2.6%) shows that an scalability gap remains to be filled regarding tuning threshold for anomaly-detection in future-researches.

As shown in latency analysis, the modelled total end-to-end response time of 231 ms per samples described by Equation (4) is within clinically-acceptable limits for asynchronous patient coaching and low-priority diagnostic interaction. The structure of global accuracy by rounds on the federated learning convergence graph also demonstrates the efficiency of communication through distributed hospital nodes since there is no variation in global model accuracy after 15 training rounds, indicating that no additional computation cycles are required. An efficient and holistic architectural approach to empowering AI-driven healthcare support for the next generation compliant with privacy regulation.

5. Conclusion

The paper presented a cybersecurity-augmented artificial intelligence agent (AI) driven conceptual model for the health care continuum from diagnosis to patient coaching. This multi-layer architecture combined transformer-based natural language processing, federated learning, differential privacy and role-based access control and real-time anomaly detection in an elegant system ripe for clinical integration. Experimentally, we found that our framework achieved diagnostic accuracy of 91.5% and F1-score of 0.91, across multiple types of adversarial attack scenarios (including prompt injection, model inversion data poisoning and STD-tester unauthorized access attempts) with an impressive overall threat detection rate of 97.3%. The convergence analysis of our federated learning confirmed that the global model consistently maintained its performance across five distributed hospital nodes and approximately 15 training rounds, validating the efficiency of our aggregation strategy. 231 ms end-to-end response latency is still the Tinderest in acceptable range for clinical threshold. These findings together add to the evidence that advanced AI conversational ability and tight cybersecurity compliance will evolve in parallel with no meaningful tradeoffs in performance. One step forward is deploying this in real hospitals, giving support for multilingual patients where necessary and optimising the detection thresholds even further to reduce higher false positive rates usually seen under complex threat environments.

References

[1] Aerts, A.; Bogdan-Martin, D. Leveraging data and AI to deliver on the promise of digital health. *Int. J. Med. Inform.* **2021**, *150*, 104456. [**Google Scholar**] [**CrossRef**] [**PubMed**]

[2] Ali, O.; Abdelbaki, W.; Shrestha, A.; Elbasi, E.; Alryalat, M.A.A.; Dwivedi, Y.K. A systematic literature review of artificial intelligence in the healthcare sector: Benefits, challenges, methodologies, and functionalities. *J. Innov. Knowl.* **2023**, *8*, 100333. [**Google Scholar**] [**CrossRef**]

[3] Li, B.H.; Hou, B.C.; Yu, W.T.; Lu, X.B.; Yang, C.W. Applications of artificial intelligence in intelligent manufacturing: A review. *Front. Inf. Technol. Electron. Eng.* **2017**, *18*, 86–96. [**Google Scholar**] [**CrossRef**]

[4] Kaplan, A.; Haenlein, M. Rulers of the world, unite! The challenges and opportunities of artificial intelligence. *Bus. Horiz.* **2020**, *63*, 37–45. [**Google Scholar**] [**CrossRef**]

[5] Chien, C.F.; Dauzere-Peres, S.; Huh, W.T.; Jang, Y.J.; Morrison, J.R. Artificial intelligence in manufacturing and logistics systems: Algorithms, applications, and case studies. *Int. J. Prod. Res.* **2020**, *58*, 2730–2731. [**Google Scholar**] [**CrossRef**]

[6] Kumar, P.; Sharma, S.K.; Dutot, V. Artificial intelligence (AI)-enabled CRM capability in healthcare: The impact on service innovation. *Int. J. Inf. Manag.* **2023**, *69*, 102598. [**Google Scholar**] [**CrossRef**]

[7] Aiken, R.M.; Epstein, R.G. Ethical guidelines for AI in education: Starting a conversation. *Int. J. Artif. Intell. Educ.* **2000**, *11*, 163–176. [**Google Scholar**]

[8] Bansal, A.; Padappayil, R.P.; Garg, C.; Singal, A.; Gupta, M.; Klein, A. Utility of artificial intelligence amidst the COVID 19 pandemic: A review. *J. Med. Syst.* **2020**, *44*, 156. [**Google Scholar**] [**CrossRef**]

[9] Chee, M.L.; Ong, M.E.H.; Siddiqui, F.J.; Zhang, Z.; Lim, S.L.; Ho, A.F.W.; Liu, N. Artificial intelligence applications for COVID-19 in intensive care and emergency settings: A systematic review. *Int. J. Environ. Res. Public Health* **2021**, *18*, 4749. [**Google Scholar**] [**CrossRef**]

[10] Minz, A.; Mahobiya, C. MR Image Classification Using ad Boost for Brain Tumor Type. In Proceedings of the IEEE 7th International Advance Computing Conference, Hyderabad, India, 5–7 January 2017; pp. 701–705. [**Google Scholar**]

[11] Schachner, T.; Keller, R.; Wangenheim, F.V. Artificial Intelligence-Based Conversational Agents for Chronic Conditions: Systematic Literature Review. *J. Med. Internet Res.* **2020**, *22*, e20701. [**Google Scholar**] [**CrossRef**]

[12] Kramer, L.L.; Ter Stal, S.; Mulder, B.; De Vet, E.; Van Velsen, L. Developing Embodied Conversational Agents for Coaching People in a Healthy Lifestyle: Scoping Review. *J. Med. Internet Res.* **2020**, *22*, e14058. [**Google Scholar**] [**CrossRef**] [**PubMed**]

[13] Ferrand, J.; Hockensmith, R.; Houghton, R.F.; Walsh-Buhi, E.R. Evaluating Smart Assistant Responses for Accuracy and Misinformation Regarding Human Papillomavirus Vaccination: Content Analysis Study. *J. Med. Internet Res.* **2020**, *22*, e19018. [**Google Scholar**] [**CrossRef**]

[14] Sezgin, E.; Militello, L.K.; Huang, Y.; Lin, S. A scoping review of patient-facing, behavioral health interventions with voice assistant technology targeting self-management and healthy lifestyle behaviors. *Transl. Behav. Med.* **2020**, *10*, 606–628. [**Google Scholar**] [**CrossRef**] [**PubMed**]

[15] Safi, Z.; Abd-Alrazaq, A.; Khalifa, M.; Househ, M. Technical Aspects of Developing Chatbots for Medical Applications: Scoping Review. *J. Med. Internet Res.* **2020**, *22*, e19127. [**Google Scholar**] [**CrossRef**] [**PubMed**]

[16] Griffin, A.C.; Xing, Z.; Khairat, S.; Wang, Y.; Bailey, S.; Arguello, J.; Chung, A.E. Conversational Agents for Chronic Disease Self-Management: A Systematic Review. In *AMIA Annual Symposium Proceedings*; American Medical Informatics Association: Bethesda, MD, USA, 2021; pp. 504–513. [**Google Scholar**]

[17] McGreevey, J.D., 3rd; Hanson, C.W., 3rd; Koppel, R. Clinical, Legal, and Ethical Aspects of Artificial Intelligence–Assisted Conversational Agents in Health Care. *JAMA J. Am. Med. Assoc.* **2020**, *324*, 552. [**Google Scholar**] [**CrossRef**] [**PubMed**]

[18] Bickmore, T.W.; Kimani, E.; Trinh, H.; Pusateri, A.; Paasche-Orlow, M.K.; Magnani, J.W. Managing Chronic Conditions with a Smartphone-based Conversational Virtual Agent. In Proceedings of the 18th International Conference on Intelligent Virtual Agents, IVA 2018, Sydney, NSW, Australia, 5–8 November 2018. [**Google Scholar**] [**CrossRef**]

[19] Pereira, J.; Díaz, Ó. Using Health Chatbots for Behavior Change: A Mapping Study. *J. Med. Syst.* **2019**, *43*, 135. [**Google Scholar**] [**CrossRef**]

[20] Greer, S.; Ramo, D.; Chang, Y.-J.; Fu, M.; Moskowitz, J.; Haritatos, J. Use of the Chatbot "Vivibot" to Deliver Positive Psychology Skills and Promote Well-Being Among Young People After Cancer Treatment: Randomized Controlled Feasibility Trial. *JMIR mHealth uHealth* **2019**, *7*, e15018. [**Google Scholar**] [**CrossRef**]